



Universiteit
Leiden
The Netherlands

Dissemination models: applications in epidemiology, rumour spreading, wealth distribution, and opinion formation

Chan, K.M.D.

Citation

Chan, K. M. D. (2026, September 16). *Dissemination models: applications in epidemiology, rumour spreading, wealth distribution, and opinion formation*. Retrieved from <https://hdl.handle.net/1887/4309746>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4309746>

Note: To cite this publication please use the final published version (if applicable).

Chapter 3

Stochastic dissemination

Chapter overview

This chapter considers a highly general dissemination model that keeps track of the stochastic evolution of the distribution of wealth over a set of agents. There are two types of events: (i) units of wealth externally arrive, and (ii) units of wealth are redistributed among the agents, while throughout Markov modulation is allowed. We derive a system of coupled differential equations describing the joint transient distribution of the agents' wealth values, which translate into *linear* differential equations when considering the corresponding means and (co-)variances. While our model uses the (economic) terminology of wealth being distributed over agents, we illustrate through a series of examples that it can be used considerably more broadly. Indeed, it also facilitates the analysis of the spread of opinions over a population (thus generalising existing opinion dynamics models), and the analysis of the dynamics of a file storage system (thus allowing the assessment of the efficacy of storage policies).

3.1 Introduction

A stylised yet highly general model that describes the spread of wealth over a population of agents could encompass the following elements. In the first place, the agents are endowed with external inflow of wealth, for instance representing their salaries, paid to them by agents outside the population. In the second place, there are transactions: agents may purchase commodities or services from other agents, thus redistributing the wealth. Thirdly, to make the model more realistic, one could impose Markov modulation on the system: all parameters involved are affected by an exogenously evolving Markovian process. This Markovian background process could for instance represent the state of the economy, alternating between economic growth and recession.

The main contributions of this chapter are the following. We develop a highly general dissemination model that keeps track of the stochastic evolution of the distribution of wealth over a set of agents, incorporating the three elements mentioned above (i.e., external inflow, redistribution and Markov modulation). A main asset of the model lies in it being broadly applicable and allowing closed-form analysis at the same time.

For the resulting model, we succeed in deriving a system of coupled differential equations that describe the joint transient probability generating function of the agents' wealth levels, jointly with the state of the Markovian background process. When focusing on the corresponding means and (co-)variances, this system takes on a more convenient form, in that it becomes a system of *linear* differential equations, thus allowing for straightforward numerical evaluation. In passing, we also consider the model's stationary behaviour, in particular establishing a stability condition.

The broad applicability of the model is illustrated through a series of examples. As suggested by our terminology, being in terms of wealth that is distributed over agents, it can be used to analyse the evolution (in time) of a wealth vector. Another example concerns the dissemination of opinions over a population, where all agents influence one another. Our modelling framework extends existing opinion dynamics models in the way we incorporate stochasticity, while also the Markovian background process is a novel element. In a last example we consider a file storage system, where files of clients are periodically copied to one or multiple central storage units. Our model can be used to assess the efficacy of policies, intended to strike a proper balance between storage cost on one hand and the risk of data loss on the other hand.

The dissemination model analysed in this chapter can be seen as a next step in a long tradition of queueing and population models. In most existing queueing network models, the dynamics are such that the number of clients per queue changes by one at a time; see e.g. the accounts in [88, 142]. A relatively small branch of the queueing literature considers queues with batch arrivals and batch services. In this respect we refer to e.g. [42, 78, 115] where product-form results are obtained. The setup that is probably closest to the one we consider in this chapter, is the one of [64], where the transaction events correspond to the network population vector undergoing a (deterministic) linear transformation. In population-process theory and epidemics there is a strong emphasis on deterministic models to describe the dynamics of the sizes of various subpopulations (e.g. age groups, infected individuals, etc.); for more background, see for instance the monograph [137]. These models' stochastic counterparts have been studied as well; a general framework has

been presented, and analysed, in [97]. Importantly, to the best of our knowledge, none of the earlier works covers our, highly general, redistribution mechanism.

This chapter is organised as follows. The model and notation are introduced in Section 3.2. Then subsequently the transient joint probability generating function (Section 3.3), the first moments (Section 3.4), and second moments — and hence also variances and covariances — (Section 3.5) are analysed; the section on first moments in addition establishes the model's stability criterion. Then there are three sections with illustrative examples, focusing on wealth redistribution (Section 3.6), on opinion dynamics (Section 3.7) and on storage sharing systems (Section 3.8). Section 3.9 concludes.

3.2 Model and notation

In our model we study the stochastic behaviour of $\mathbf{M}(t) \equiv (M_1(t), \dots, M_I(t))$, where $M_i(t)$ denotes the “wealth” of agent i at time t , with $i = 1, \dots, I$, for some $I \in \mathbb{N}$; in this chapter we follow the convention that bold symbols are used to represent vectors. We recall that, as pointed out in the introduction, “wealth” is to be interpreted in the broad sense; as we will extensively argue, the setup considered can also be used e.g. in the context of opinion spreading dynamics, or the context of file storage systems.

To make our model as rich as possible, we let its dynamics be affected by an autonomously evolving Markovian background (or regime-switching) process. In the (economic) context of wealth being spread over a population of individuals, the background process could reflect the state of the economy (e.g. alternating between economic peaks and periods of recession). Let this regime-switching process be modelled by the continuous-time Markov process $(X(t))_{t \geq 0}$ on the state space $\{1, \dots, d\}$, for some $d \in \mathbb{N}$. This process, which is assumed to be irreducible, is governed by the transition rate matrix $Q = \{q_{ij}\}_{i,j=1}^d$ (with all non-diagonal elements being non-negative and row sums equal to 0), so that [9, 124]

$$\mathbb{P}(X(t) = \ell \mid X(0) = k) = (e^{Qt})_{k,\ell}.$$

We proceed by describing the dynamics of the wealth process, given the background process is in state $k \in \{1, \dots, d\}$. We distinguish two types of events.

- In the first place there are “external arrivals” of wealth. Concretely, for $j \in \{1, \dots, J\}$ with $J \in \mathbb{N}$, at Poisson epochs with rate $\lambda_{jk} > 0$ these external arrivals occur, leading to an increase of the wealth of all agents $i \in S_j \subseteq \{1, \dots, I\}$ by one unit.
- In the second place there are “shocks”, arriving to the system according to a Poisson process with rate $\gamma_k > 0$. At such a shock, “transactions” take place, which concretely means that each of the $M_i(t)$ wealth units of agent i contributes $W_{ijk} \in \mathbb{N}_0$ wealth units to agent j .

The precise mechanism is described more formally as follows. Supposing a shock happens at time $t > 0$, then the number of wealth units after the shock at agent j is, conditional on $\mathbf{M}(t-) = (m_1, \dots, m_I)^\top$ being the wealth vector just prior to time t , given by

$$\sum_{i=1}^I \sum_{n=1}^{m_i} W_{ijkn},$$

with $(W_{ijkn})_{n \in \mathbb{N}}$ denoting a sequence of independent and identically distributed random variables, all of them distributed as the discrete, non-negative random variable W_{ijk} . The random variables W_{ijk} (with $i = 1, \dots, I$ and $k = 1, \dots, d$) are assumed independent; importantly, throughout we *do* allow dependence in j . We define, for $\mathbf{z}$ such that $\max\{|z_1|, \dots, |z_I|\} \leq 1$, the associated probability generating function (PGF) by

$$g_{ik}(\mathbf{z}) = \mathbb{E} \left[\prod_{j=1}^I z_j^{W_{ijk}} \right].$$

Our primary aim is to establish a unique characterisation of the transient wealth vector $\mathbf{M}(t)$, jointly with the state of the background process $X(t)$. To this end we work with the corresponding multivariate time-dependent joint PGF. Concretely, our analysis aims at identifying the following key object of study, for $\mathbf{z}$ such that $\max\{|z_1|, \dots, |z_I|\} \leq 1$:

$$f_k(\mathbf{z}, t) := \mathbb{E} \left[\prod_{j=1}^I z_j^{M_j(t)} \mathbb{1}_{\{X(t)=k\}} \right], \quad (3.1)$$

which uniquely defines the distribution of $(\mathbf{M}(t), X(t)) \in \mathbb{N}^I \times \{1, \dots, d\}$. In the remainder of the chapter we have used, as systematically as possible, $i, j \in \{1, \dots, I\}$ to denote agent types, $k, \ell \in \{1, \dots, d\}$ states of the background process, and $n \in \mathbb{N}$ the index of the sampled random variable.

3.3 Derivation of the differential equation for the joint PGF

The main objective of this section is to establish a system of coupled differential equations (in t) for the time-dependent joint PGFs $f_k(\mathbf{z}, t)$, as defined in (3.1). We do so relying on a standard argumentation: we relate their values at time $t + \Delta t$ to their values at time t , with the aim to set up a system of differential equations. To this end, the underlying idea is to distinguish the three types of events that can occur in an interval of length Δt : a transition of the background process, external arrivals, and shocks (and, evidently, there is in addition the event that none of these three types of events occurs). Following this line of reasoning, we obtain for the time-dependent joint PGF at time $t + \Delta t$ that, as $\Delta t \downarrow 0$,

$$\begin{aligned} f_k(\mathbf{z}, t + \Delta t) &= \sum_{\ell=1, \ell \neq k}^d q_{\ell k} \Delta t \mathbb{E} \left[\prod_{j=1}^I z_j^{M_j(t)} \mathbb{1}_{\{X(t)=\ell\}} \right] + \\ &\quad \sum_{j=1}^J \lambda_{jk} \Delta t \left(\prod_{i \in S_j} z_i \right) \mathbb{E} \left[\prod_{i=1}^I z_i^{M_i(t)} \mathbb{1}_{\{X(t)=k\}} \right] + \\ &\quad \gamma_k \Delta t \sum_{\mathbf{m} \in \mathbb{N}^I} \mathbb{E} \left[\prod_{j=1}^I z_j^{M_j(t+\Delta t)} \mathbb{1}_{\{X(t+\Delta t)=k\}} \mid \mathbf{M}(t) = \mathbf{m}, \mathcal{E}_k(t) \right] \mathbb{P}(\mathbf{M}(t) = \mathbf{m}) + \\ &\quad \left(1 - \sum_{\ell=1, \ell \neq k}^d q_{\ell k} \Delta t - \sum_{j=1}^J \lambda_{jk} \Delta t - \gamma_k \Delta t \right) f_k(\mathbf{z}, t) + o(\Delta t), \end{aligned} \quad (3.2)$$

with $\mathcal{E}_k(t)$ denoting the event of a shock between times t and $t + \Delta t$ (evidently, while the background state is k). Above we used the Landau notation $o(\Delta t)$: stating that a function $F(\cdot)$ is $o(\Delta t)$ as $\Delta t \downarrow 0$ means that $F(\Delta t)/\Delta t \rightarrow 0$ as $\Delta t \downarrow 0$.

The right-hand side of Equality (3.2) can be interpreted and rewritten as follows.

- The first term, which considers the scenario that the background process was in a state $\ell \neq k$ at time t and makes a transition to k between t and Δt , equals by definition

$$\sum_{\ell=1, \ell \neq k}^d q_{\ell k} \Delta t f_{\ell}(\mathbf{z}, t).$$

- The second term represents the contributions of the external arrivals: if it is of the j -th type, then it increases (by 1) the wealth values of all agents i such that $i \in S_j$. It reads

$$\sum_{j=1}^J \lambda_{jk} \Delta t \left(\prod_{i \in S_j} z_i \right) f_k(\mathbf{z}, t).$$

- The third term describes the effect of the shocks. The claim is that we can express it in terms of the PGF $f_k(\mathbf{z}, t)$, but not evaluated in the argument $\mathbf{z}$ but rather in a different argument $\mathbf{h}_k(\mathbf{z})$. Indeed, defining

$$\mathbf{h}_k(\mathbf{z}) := (g_{1k}(\mathbf{z}), \dots, g_{Ik}(\mathbf{z})),$$

again up to Δt -terms, due to the shock that occurs between times t and $t + \Delta t$,

$$\begin{aligned} & \sum_{\mathbf{m} \in \mathbb{N}^I} \mathbb{E} \left[\prod_{j=1}^I z_j^{M_j(t+\Delta t)} \mathbb{1}_{\{X(t+\Delta t)=k\}} \mid \mathbf{M}(t) = \mathbf{m}, \mathcal{E}_k(t) \right] \mathbb{P}(\mathbf{M}(t) = \mathbf{m}) \\ &= \sum_{\mathbf{m} \in \mathbb{N}^I} \prod_{j=1}^I (g_{jk}(\mathbf{z}))^{m_j} \mathbb{P}(\mathbf{M}(t) = \mathbf{m}, X(t) = k) \\ &= \mathbb{E} \left[\prod_{j=1}^I g_{jk}(\mathbf{z})^{M_j(t)} \mathbb{1}_{\{X(t)=k\}} \right] = f_k(\mathbf{h}_k(\mathbf{z}), t). \end{aligned}$$

- The fourth term corresponds to the scenario of no transition of the background process, no external arrivals and no shocks, leaving the wealth process unchanged.

Observe that we have succeeded in expressing $f_k(\mathbf{z}, t + \Delta t)$ in terms of quantities of the same type, as well as quantities of the type $f_{\ell}(\boldsymbol{\varphi}(\mathbf{z}), t)$ for known functions $\boldsymbol{\varphi}(\mathbf{z}) : [-1, 1]^I \rightarrow [-1, 1]^I$. The next step is to subtract $f_k(\mathbf{z}, t)$ from both sides of the equation, divide by Δt , and let $\Delta t \downarrow 0$, so as to obtain a system of differential equations in t . Indeed, we obtain, for $t \geq 0$, the following result, where we have used that the row sums of Q are equal to 0.

Proposition 3.3.1. *For any $t \geq 0$, $f_k(\mathbf{z}, t)$ satisfies the system of differential equations*

$$\begin{aligned} \frac{\partial}{\partial t} f_k(\mathbf{z}, t) &= \sum_{\ell=1}^d q_{\ell k} f_{\ell}(\mathbf{z}, t) + \\ & \sum_{j=1}^J \lambda_{jk} \left(\prod_{i \in S_j} z_i - 1 \right) f_k(\mathbf{z}, t) + \gamma_k (f_k(\mathbf{h}_k(\mathbf{z}), t) - f_k(\mathbf{z}, t)), \end{aligned} \tag{3.3}$$

If $\mathbf{M}(0) = \mathbf{m}_0$ and $X(0) = k_0$, then the initial condition is

$$f_k(\mathbf{z}, 0) = \mathbb{1}_{\{k=k_0\}} \prod_{i=1}^I z_i^{m_{0,i}}.$$

Using the obvious property that $\mathbf{h}_k(\mathbf{1}) = \mathbf{1}$, it is readily seen from this proposition that $f_k(\mathbf{1}, t) = \mathbb{P}(X(t) = k)$, as it should.

The system of differential equations (3.3) can be solved numerically. Importantly, the system is *not* linear; observe that in one term on the right-hand side of (3.3) the PGF has the argument $\mathbf{h}_k(\mathbf{z})$ rather than $\mathbf{z}$. As we will notice in the next sections, however, when considering the computation of time-dependent moments (rather than the full time-dependent PGF) we *do* obtain a reduction to systems of differential equations that *are* linear. These can be solved at relatively low numerical effort using standard computational software (e.g. `deSolve` in R, `dsolve` in Matlab). Alternatively, as in our case the differential equations are linear, one can use an implementation of the matrix exponential (`expm` in both R and Matlab). In the next two sections we subsequently concentrate on the evaluation of the first and second moments.

Remark 3.3.2. If one is interested in moments only, then one could in principle also set up systems of differential equations just for those moments, rather than setting up a system of differential equations for the PGF and then perform a differentiation (as will be done in the next sections); cf. for instance the procedure followed in [100]. We however prefer to present the system of coupled ordinary differential equations of Proposition 3.1, as these characterise the *full* distribution of $\mathbf{M}(t)$ (jointly with the state of the background process $X(t)$). These coupled differential equations do not allow an analytical solution, but can be solved numerically relying on standard computational software packages (use e.g. the function `deSolve` in R, or `dsolve` in Matlab). It is also noted that repeated differentiation with respect to entries of $\mathbf{z}$ and inserting $\mathbf{z} = \mathbf{0}$ (i.e., *not* $\mathbf{z} = \mathbf{1}$) yields the joint probability distribution; for instance, for a single-dimensional random variable X attaining non-negative integer values, we have

$$\mathbb{P}(X = k) = \frac{1}{k!} \left. \frac{d^k}{dz^k} \mathbb{E} z^X \right|_{z=0},$$

for $k \in \mathbb{N}$.

3.4 Derivation of first moments, stability

In principle all moments of the components of $\mathbf{M}(t)$, as well as all mixed moments, can be derived from the differential equations (3.3) by repeated differentiation and plugging in $\mathbf{z} = \mathbf{1}$. This typically leads to a recursion from which the joint moments of order n are expressed in terms of their counterparts of order $0, 1, \dots, n-1$; cf. [25, 101, 149]. In this section we discuss this widely applied procedure to determine the time-dependent first moments. In the next section we then move to second moments, but in principle this procedure extends to moments of any order. As we point out in the present section, knowledge of the first moments also provides us with a criterion under which the model has a stable stationary version (i.e., does not explode as $t \rightarrow \infty$).

3.4.1 Differential equations for first moments

In this subsection we derive a system of linear differential equations that characterise the expectation of $\mathbf{M}(t)$. With $w_{ijk} := \mathbb{E}W_{ijk}$ and $m_{ik}(t) := \mathbb{E}[M_i(t) \mathbb{1}_{\{X(t)=k\}}]$, differentiating (3.3) to z_i and inserting $\mathbf{z} = \mathbf{1}$ yields the following system of coupled ordinary differential equations:

$$m'_{ik}(t) = \sum_{\ell=1}^d q_{\ell k} m_{i\ell}(t) + \sum_{j:i \in S_j} \lambda_{jk} \pi_k(t) + \gamma_k \left(\sum_{j=1}^I w_{jik} m_{jk}(t) - m_{ik}(t) \right),$$

with $\pi_k(t) = \mathbb{P}(X(t) = k)$. Here we have used the standard differentiation rule for compositions of functions with vector-valued arguments, i.e.,

$$\frac{\partial f_k(\mathbf{h}_k(\mathbf{z}), t)}{\partial z_i} = \sum_{j=1}^I \frac{\partial f_k(\mathbf{x}, t)}{\partial x_j} \Big|_{\mathbf{x}=\mathbf{h}_k(\mathbf{z})} \frac{\partial (h_k(\mathbf{z}))_j}{\partial z_i},$$

and

$$\frac{\partial (h_k(\mathbf{z}))_j}{\partial z_i} \Big|_{\mathbf{z}=\mathbf{1}} = w_{jik}.$$

The next step is to compactly write the above system of differential equations in matrix-vector form. We let $\mathbf{m}(t) \in \mathbb{R}^{dI}$ denote the stacked vector that results from the I vectors $\mathbf{m}_i(t) \equiv (m_{i1}(t), \dots, m_{id}(t))^{\top}$, where $\boldsymbol{\pi}(t) \in \mathbb{R}^{dI}$ is the stacked vector that results from the I (identical) vectors $\boldsymbol{\pi}_i(t) \equiv (\pi_1(t), \dots, \pi_d(t))^{\top}$, i.e.,

$$\mathbf{m}(t) := \begin{pmatrix} \mathbf{m}_1(t) \\ \vdots \\ \mathbf{m}_I(t) \end{pmatrix}, \quad \boldsymbol{\pi}(t) := \begin{pmatrix} \boldsymbol{\pi}_1(t) \\ \vdots \\ \boldsymbol{\pi}_I(t) \end{pmatrix}.$$

In addition, $G_{ji} := \text{diag}\{\gamma_1 w_{ji1}, \dots, \gamma_d w_{jid}\} - \text{diag}\{\gamma_1, \dots, \gamma_d\} \mathbb{1}_{\{i=j\}}$. We then define

$$A := \begin{pmatrix} Q^{\top} + G_{11} & G_{21} & G_{31} & \cdots & G_{I1} \\ G_{12} & Q^{\top} + G_{22} & G_{32} & \cdots & G_{I2} \\ G_{13} & G_{23} & Q^{\top} + G_{33} & \cdots & G_{I3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ G_{1I} & G_{2I} & G_{3I} & \cdots & Q^{\top} + G_{II} \end{pmatrix},$$

and, with $\Lambda_i := \text{diag}\{\bar{\lambda}_{i1}, \dots, \bar{\lambda}_{id}\}$ and $\bar{\lambda}_{ik} := \sum_{j:i \in S_j} \lambda_{jk}$,

$$\Lambda := \begin{pmatrix} \Lambda_1 & 0 & \cdots & 0 \\ 0 & \Lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \Lambda_I \end{pmatrix}.$$

We thus end up with a system of dI coupled linear differential equations, as given in the following proposition.

Proposition 3.4.1. *For any $t \geq 0$, $\mathbf{m}(t)$ satisfies the system of linear differential equations*

$$\mathbf{m}'(t) = A \mathbf{m}(t) + \Lambda \boldsymbol{\pi}(t).$$

If $\mathbf{M}(0) = \mathbf{m}_0$ and $X(0) = k_0$, then the initial condition is

$$m_k(0) = \mathbb{1}_{\{k=k_0\}} m_{0,k}.$$

This non-homogeneous system of linear differential equations can be solved in the standard manner. In the first place, the vector $\boldsymbol{\pi}(t)$, corresponding to the transient state probabilities of the background process $X(t)$, satisfies the differential equation

$$\boldsymbol{\pi}'(t) = (\mathbb{I}_I \otimes Q^\top) \boldsymbol{\pi}(t),$$

with $\otimes$ being the usual notation for the Kronecker product and $\mathbb{I}_I$ an identity matrix of dimension I . This means, with $\bar{Q}$ denoting the $(dI \times dI)$ -dimensional matrix $\mathbb{I}_I \otimes Q^\top$, that $\boldsymbol{\pi}(t) = e^{\bar{Q}t} \boldsymbol{\pi}(0)$, or equivalently,

$$\pi_j(t) = \sum_{i=1}^d \mathbb{P}(X(0) = i) (e^{\bar{Q}t})_{i,j}.$$

In the second place, the solution for $\mathbf{m}(t)$ can be written in terms of matrix exponentials, as follows:

$$\mathbf{m}(t) = e^{At} \mathbf{m}(0) + \int_0^t e^{A(t-s)} \Lambda \boldsymbol{\pi}(s) ds.$$

We thus obtain the following result.

Proposition 3.4.2. *For any $t \geq 0$,*

$$\mathbf{m}(t) = e^{At} \mathbf{m}(0) + \int_0^t e^{A(t-s)} \Lambda e^{\bar{Q}s} \boldsymbol{\pi}(0) ds.$$

3.4.2 Stability condition

Let $\boldsymbol{\pi} := \boldsymbol{\pi}(\infty)$ be the unique solution of $\boldsymbol{\pi}Q = \mathbf{0}$ such that its entries sum to 1, i.e., the stationary distribution of $X(t)$. In case the underlying model is stable, the above results directly imply that the steady-state mean vector $\mathbf{m}$ can be written in terms of the steady-state probabilities $\boldsymbol{\pi}$, as follows:

$$\mathbf{m} = -A^{-1} \Lambda \boldsymbol{\pi}.$$

The formal stability condition is given in the next statement. We define by ω the eigenvalue of A with largest real part, i.e., the spectral abscissa of A .

Proposition 3.4.3. *The Markov chain $(\mathbf{M}(t), X(t))_{t \geq 0}$ is ergodic if $\omega < 0$.*

Proof. The proof step-by-step mimics the one of [64, Prop. 3]; we therefore restrict ourselves to sketching its main steps. The underlying idea is to establish ergodicity of the skeleton Markov chain $(\mathbf{M}(n\Delta), X(n\Delta))_{n \in \mathbb{N}}$ if $\omega < 0$, for some $\Delta > 0$, where it should be noted

that if the skeleton Markov chain is ergodic for some $\Delta > 0$, then so is $(\mathbf{M}(t), X(t))_{t \geq 0}$ (observing that the mean recurrence time for any state of the skeleton chain is an upper bound for the mean recurrence time of the original process). Then, by [9, Prop. I.5.3], a sufficient condition for ergodicity can be phrased in terms of

$$\mathbb{E}[\|\mathbf{M}(\Delta)\|_1 \mid \mathbf{M}(0) = \mathbf{m}_0, X(0) = k_0] - \|\mathbf{m}_0\|_1 < -\varepsilon,$$

for some $\varepsilon > 0$, and all $\mathbf{m}_0 \in \mathbb{N}^I$ and $k_0 \in \{1, \dots, d\}$. Informally, this criterion entails that the process' drift is negative, bounded away from zero; cf. Foster's criterion [27, 69]. Such a bound can be achieved under $\omega < 0$, with the precise same argumentation as the one used in the proof of [64, Prop. 3], where we rely on [20, Prop. 11.18] to find the required bound on the norm of the matrix exponential. $\square$

In case the stability condition $\omega < 0$ is not fulfilled, the components of $\mathbf{m}(t)$ typically grow in a fixed proportion; we return to this issue in Section 3.7.

3.4.3 Special cases

In this subsection we provide more explicit results for three special cases: (1) a single fully homogeneous population, (2) a single distinct agent (a “leader”) with $I - 1$ homogeneous other agents (“followers”), (3) two, internally homogeneous, interacting subpopulations.

Homogeneous population.

We denote by $m_k(t)$ the mean wealth of an arbitrary agent, jointly with the event that the background process is in state k . As a consequence of the fact that we “start symmetrically”, i.e., the initial wealth of all agents is the same, at any point in time the mean wealth of the individual agents coincides. In case we do not “start symmetrically” (for instance with two possible values of the initial wealth), the computations can still be performed, but become less clean.

In the variant that we consider, we let $J = I$, and we take $S_j = \{j\}$ and $\lambda_{jk} = \lambda_k$, for $j = 1, \dots, I$. (Here we note that one can construct other fully symmetric external arrival processes, for instance by letting $J = 1$ and $S_1 = \{1, \dots, I\}$. Such alternative symmetric variants can be dealt with analogously.) The generic random variable W_{ijk} now depends on the background state k only (i.e., not on the indices i and j that indicate the agents); we let w_k be its expected value.

It is directly verified that our earlier results now yield

$$m'_k(t) = \sum_{\ell=1}^d q_{\ell k} m_\ell(t) + \lambda_k \pi_k(t) + \gamma_k (I w_k - 1) m_k(t),$$

or, in matrix-vector notation,

$$\mathbf{m}'(t) = A \mathbf{m}(t) + \Lambda \boldsymbol{\pi}(t),$$

where $A := Q^\top + \text{diag}\{\gamma_1(I w_1 - 1), \dots, \gamma_d(I w_d - 1)\}$ and $\Lambda := \text{diag}\{\lambda_1, \dots, \lambda_d\}$. The k th entry of $\mathbf{m}$ now expresses the agents' mean wealth when the system is in state k . In

steady state, we obtain that the mean wealth vector equals $-A^{-1}\Lambda\boldsymbol{\pi}$, with $\boldsymbol{\pi}$ as defined before, provided that the stability condition is fulfilled (i.e., that the spectral abscissa of A is negative).

Leader and homogeneous followers.

In this model, there is a single leader and $I - 1$ homogeneous followers. We “start symmetrically”, i.e., all followers have the same initial wealth. We let $m_{L,k}(t)$ be the mean wealth of the leader at time t , and $m_{F,k}(t)$ the mean wealth of an arbitrary follower at time t , both jointly with the event that the background process is in state k .

As we did in the case of a homogeneous population, we take $J = I$ with $S_j = \{j\}$. We let the external arrival rate of the leader be $\lambda_{L,k}$, and of the followers $\lambda_{F,k}$, both corresponding to the background process being in state k . The wealth redistribution, as taking place at the “shocks”, corresponds to the means (in self-evident notation) $w_{LL,k}$, $w_{LF,k}$, $w_{FL,k}$, and $w_{FF,k}$, again for the background process in state k . We thus obtain

$$\begin{aligned} m'_{L,k}(t) &= \sum_{\ell=1}^d q_{\ell k} m_{L,\ell}(t) + \lambda_{L,k} \pi_k(t) + \gamma_k((w_{LL,k} - 1)m_{L,k}(t) + (I - 1)w_{FL,k}m_{F,k}(t)), \\ m'_{F,k}(t) &= \sum_{\ell=1}^d q_{\ell k} m_{F,\ell}(t) + \lambda_{F,k} \pi_k(t) + \gamma_k((w_{FF,k}(I - 1) - 1)m_{F,k}(t) + w_{LF,k}m_{L,k}(t)). \end{aligned}$$

The model further simplifies if we consider the setting without modulation. In self-evident notation, we obtain $\mathbf{m}'(t) = \gamma(\bar{A} - \mathbb{I}_2)\mathbf{m}(t) + \boldsymbol{\lambda}$, where $\mathbb{I}_2$ denotes a 2-dimensional identity matrix,

$$\mathbf{m}(t) = \begin{pmatrix} m_L(t) \\ m_F(t) \end{pmatrix}, \quad \bar{A} := \begin{pmatrix} w_{LL} & (I - 1)w_{FL} \\ w_{LF} & (I - 1)w_{FF} \end{pmatrix}, \quad \boldsymbol{\lambda} = \begin{pmatrix} \lambda_L \\ \lambda_F \end{pmatrix}.$$

Two internally homogeneous interacting subpopulations

We now consider a generalisation of the situation with a leader and homogeneous followers, viz. the situation of I_A agents of subpopulation A and $I_B := I - I_A$ agents of subpopulation B. We let $m_{A,k}(t)$ (respectively $m_{B,k}(t)$) denote the mean wealth of an arbitrary agent from subpopulation A (respectively subpopulation B) at time t ; the agents within each of the two subpopulations “start symmetrically”. The arrival rates $\lambda_{A,k}$ and $\lambda_{B,k}$ are defined in the evident manner, and so are the means $w_{AA,k}$, $w_{AB,k}$, $w_{BA,k}$, and $w_{BB,k}$. Using the same reasoning as above, we obtain

$$\begin{aligned} m'_{A,k}(t) &= \sum_{\ell=1}^d q_{\ell k} m_{A,\ell}(t) + \lambda_{A,k} \pi_k(t) + \gamma_k((I_A w_{AA,k} - 1)m_{A,k}(t) + I_B w_{BA,k}m_{B,k}(t)), \\ m'_{B,k}(t) &= \sum_{\ell=1}^d q_{\ell k} m_{B,\ell}(t) + \lambda_{B,k} \pi_k(t) + \gamma_k((I_B w_{BB,k} - 1)m_{B,k}(t) + I_A w_{AB,k}m_{A,k}(t)). \end{aligned}$$

We get a further simplification in case that there is no modulation. In self-evident notation, we again obtain $\mathbf{m}'(t) = \gamma(\bar{A} - \mathbb{I}_2)\mathbf{m}(t) + \boldsymbol{\lambda}$, but now with

$$\mathbf{m}(t) = \begin{pmatrix} m_A(t) \\ m_B(t) \end{pmatrix}, \quad \bar{A} := \begin{pmatrix} I_A w_{AA} & I_B w_{BA} \\ I_A w_{AB} & I_B w_{BB} \end{pmatrix}, \quad \boldsymbol{\lambda} = \begin{pmatrix} \lambda_A \\ \lambda_B \end{pmatrix}.$$

3.5 Derivation of second moments

In this section we focus on characterising the second moments pertaining to the vector $\mathbf{M}(t)$. The techniques relied upon resemble those used in the previous section to find the first moments. In particular, the solution again amounts to solving a system of linear differential equations.

3.5.1 Differential equations for second moments

Concretely, our aim is to provide recipes to evaluate the reduced second moments of $M_i(t)$, i.e.,

$$v_{iik}(t) := \mathbb{E}[M_i(t)(M_i(t) - 1)\mathbb{1}_{\{X(t)=k\}}] = \frac{\partial^2 f_k(\mathbf{z}, t)}{\partial z_i^2},$$

as well as the mixed second moments of $M_i(t)$ and $M_{i'}(t)$, i.e., for $i \neq i'$,

$$v_{ii'k}(t) := \mathbb{E}[M_i(t) M_{i'}(t)\mathbb{1}_{\{X(t)=k\}}] = \frac{\partial^2 f_k(\mathbf{z}, t)}{\partial z_i \partial z_{i'}}.$$

Again the idea is to set up a system of coupled linear differential equations. In these differential equations now both the transient state probabilities $\pi_k(t)$ and the transient first moments $m_{ik}(t)$ feature. With these objects at our disposal (recalling in particular that an expression for $m_{ik}(t)$ was identified in the previous section), we can determine the corresponding variances and covariances in the evident manner.

In the derivation, we use the identity, for $i, i' = 1, \dots, I$,

$$\begin{aligned} \frac{\partial^2 f_k(\mathbf{h}_k(\mathbf{z}), t)}{\partial z_i \partial z_{i'}} &= \frac{\partial}{\partial z_i} \left(\sum_{j'=1}^I \frac{\partial f_k(\mathbf{x}, t)}{\partial x_{j'}} \Big|_{\mathbf{x}=\mathbf{h}_k(\mathbf{z})} \frac{\partial (h_k(\mathbf{z}))_{j'}}{\partial z_{i'}} \right) \\ &= \sum_{j=1}^I \sum_{j'=1}^I \frac{\partial^2 f_k(\mathbf{x}, t)}{\partial x_j \partial x_{j'}} \Big|_{\mathbf{x}=\mathbf{h}_k(\mathbf{z})} \frac{\partial (h_k(\mathbf{z}))_j}{\partial z_i} \frac{\partial (h_k(\mathbf{z}))_{j'}}{\partial z_{i'}} \\ &\quad + \sum_{j'=1}^I \frac{\partial f_k(\mathbf{x}, t)}{\partial x_{j'}} \Big|_{\mathbf{x}=\mathbf{h}_k(\mathbf{z})} \frac{\partial^2 (h_k(\mathbf{z}))_{j'}}{\partial z_i \partial z_{i'}}, \end{aligned}$$

which is based on standard differentiation rules. We thus obtain, by differentiating (3.3) with respect to z_i and $z_{i'}$ and inserting $\mathbf{z} = \mathbf{1}$,

$$\begin{aligned} v'_{ii'k}(t) &= \sum_{\ell=1}^d q_{\ell k} v_{ii'\ell}(t) + \mathbb{1}_{\{i \neq i'\}} \sum_{j:i, i' \in S_j} \lambda_{jk} \pi_k(t) + \sum_{j:i \in S_j} \lambda_{jk} m_{i'k}(t) + \sum_{j:i' \in S_j} \lambda_{jk} m_{ik}(t) + \\ &\quad \gamma_k \left(\sum_{j=1}^I \sum_{j'=1}^I v_{jj'k}(t) w_{jik} w_{j'i'k} + \sum_{j=1}^I m_{jk}(t) w_{jii'k}^{(2)} - v_{ii'k}(t) \right); \end{aligned}$$

here

$$w_{jii'k}^{(2)} := \frac{\partial^2 (h_k(\mathbf{z}))_j}{\partial z_i \partial z_{i'}} \Big|_{\mathbf{z}=\mathbf{1}},$$

which equals $\mathbb{E}[W_{jik}(W_{jik} - 1)]$ if $i = i'$ and $\mathbb{E}[W_{jik}W_{j'i'k}]$ otherwise.

3.5.2 Special case: two subpopulations

We consider the situation of Section 3.4.3 and derive the differential equations for the (reduced) second moments in case there is no modulation. Five quantities are to be determined:

- $v_{AA}(t) :=$ reduced 2nd moment of arbitrary agent in population A,
- $v_{BB}(t) :=$ reduced 2nd moment of arbitrary agent in population B,
- $v_{AA'}(t) :=$ mixed 2nd moment of two arbitrary distinct agents in population A,
- $v_{BB'}(t) :=$ mixed 2nd moment of two arbitrary distinct agents in population B,
- $v_{AB}(t) :=$ mixed 2nd moment of two arbitrary agents in populations A and B,

with all quantities on the right-hand side being evaluated at time $t \geq 0$. The vector $\mathbf{v}(t) \in \mathbb{R}_+^5$ consists of the above five entries. We can write, with $\mathbb{I}_5$ denoting a 5-dimensional identity matrix,

$$\mathbf{v}'(t) = \gamma(\bar{A} - \mathbb{I}_5) \mathbf{v}(t) + A_m \mathbf{m}(t),$$

for a suitably chosen (5×5) -matrix $\bar{A}$ and a suitably chosen (5×2) -matrix A_m . The matrix $\bar{A}$ is given by, with $J_x := I_x(I_x - 1)$ for $x \in \{A, B\}$,

$$\bar{A} = \begin{pmatrix} I_A(w_{AA})^2 & I_B(w_{BA})^2 & J_A(w_{AA})^2 & J_B(w_{BA})^2 & 2I_A I_B w_{AA} w_{BA} \\ I_A(w_{AB})^2 & I_B(w_{BB})^2 & J_A(w_{AB})^2 & J_B(w_{BB})^2 & 2I_A I_B w_{AB} w_{BB} \\ I_A(w_{AA})^2 & I_B(w_{BA})^2 & J_A(w_{AA})^2 & J_B(w_{BA})^2 & 2I_A I_B w_{AA} w_{BA} \\ I_A(w_{AB})^2 & I_B(w_{BB})^2 & J_A(w_{AB})^2 & J_B(w_{BB})^2 & 2I_A I_B w_{AB} w_{BB} \\ I_A w_{AA} w_{AB} & I_B w_{BA} w_{BB} & J_A w_{AA} w_{AB} & J_B w_{BA} w_{BB} & I_A I_B (w_{AA} w_{BB} + w_{AB} w_{BA}) \end{pmatrix}$$

and

$$A_m = \begin{pmatrix} 2\lambda_A + \gamma I_A w_{AAA}^{(2)} & \gamma I_B w_{BAA}^{(2)} \\ \gamma I_A w_{ABB}^{(2)} & 2\lambda_B + \gamma I_B w_{BBB}^{(2)} \\ 2\lambda_A + \gamma I_A w_{AAA'}^{(2)} & \gamma I_B w_{BAA'}^{(2)} \\ \gamma I_A w_{ABB'}^{(2)} & 2\lambda_B + \gamma I_B w_{BBB'}^{(2)} \\ \lambda_A + \gamma I_A w_{AAB}^{(2)} & \lambda_B + \gamma I_B w_{BAB}^{(2)} \end{pmatrix}.$$

3.6 Application 1: wealth redistribution

The model that we consider in this chapter can be interpreted as a simple formalism describing an economy, providing insight into the stochastic evolution of the wealth of each of the individual agents. Indeed, the set I could correspond to the agents of the economic system under study, which is fed by external inflow and in which at random times wealth redistribution occurs. In this section we present an example of such an economic system, with a population consisting of one agent (the “leader”) obtaining income from outside the system, and the other $I - 1$ agents obtaining their income from the leader (the “followers”). The Markovian background process records the state of the economy, in that it alternates between periods of economic growth and periods of recession. Our model allows us to quantify the distribution of the fraction of followers whose income drops below a critical threshold. In this way, we can get insight into the phenomenon of “poverty trap”, i.e., persistent poverty for the followers (to be interpreted as the segment of the population

that does not own resources, and whose income strongly depends on payments from the leader). The model considered is formally described as follows.

- The background process $X(t)$ has two states, i.e., economic growth (corresponding to state 1) and recession (corresponding to state 2).

- Regarding the income rates λ_j , we consider the situation that $S_j = \{j\}$, for $j = 1, \dots, I$. Only agent 1 (the leader) has external income: we let $\lambda_{j1} = \lambda_1$ and $\lambda_{j2} = \lambda_2$ for rates λ_1 and λ_2 such that $0 < \lambda_2 < \lambda_1$ (i.e., the leader has a higher income during periods of economic growth). The other agents (i.e., the followers) do not have any external income: $\lambda_{jk} = 0$ for $j = 2, \dots, I$.

- Wealth distribution takes place at a Poisson rate γ_k , with k the state of the background process. The vectors $(W_{11k}, \dots, W_{1Ik})$, for $k = 1, 2$, are multinomially distributed with parameters 1 and $(p_k, r_k, \dots, r_k)$. Here $r_k \leq (1 - p_k)/(I - 1)$, meaning that with probability $1 - p_k - (I - 1)r_k \in [0, 1]$ the wealth unit of agent 1 leaves the economy. Typically one expects $p_2 > p_1$, as in periods of recession the leader will be inclined to save a larger fraction of their wealth. We let the vectors $(W_{j1k}, \dots, W_{jIk})$, for $k = 1, 2$ and $j = 2, \dots, I$, be multinomially distributed with parameters 1 and $(0, s_k, \dots, s_k)$. Here $s_k \leq 1/(I - 1)$, meaning that with probability $1 - (I - 1)s_k \in [0, 1]$ the wealth unit of client j leaves the economy.

In this setup intentionally various symmetries are assumed, so as to keep the model as low-dimensional as possible. Evidently, more general variants can be analysed as well. For example, we could consider instances such that, for $j = 2, \dots, I$ and $j' \neq j$, the distribution of W_{jjk} differs from the distribution of $W_{jj'k}$.

3.6.1 Means and variances

It is not hard to verify that

$$g_{1k}(\mathbf{z}) = z_1^{p_k} \left(\prod_{j=2}^I z_j \right)^{r_k}, \quad g_{jk}(\mathbf{z}) = \left(\prod_{j=2}^I z_j \right)^{s_k},$$

thus also defining $h_k(\mathbf{z})$. For ease, assume that the background process is in stationarity at time 0. We thus obtain the following differential equations for the transient means (in self-evident notation):

$$\begin{aligned} m'_{L,k}(t) &= \sum_{\ell=1}^2 q_{\ell k} m_{L,\ell}(t) + \lambda_k \pi_k + \gamma_k (p_k - 1) m_{L,k}(t), \\ m'_{F,k}(t) &= \sum_{\ell=1}^2 q_{\ell k} m_{F,\ell}(t) + \gamma_k (((I - 1)s_k - 1) m_{F,k}(t) + r_k m_{L,k}(t)), \end{aligned}$$

using the results found in Section 3.4.3. The transient means as described by the differential equations above indeed nicely match simulation results as shown in Figure 3.1. Focusing

on stationarity, we obtain that

$$\begin{pmatrix} m_{L,1} \\ m_{L,2} \end{pmatrix} = - \begin{pmatrix} -q_1 + \gamma_1(p_1 - 1) & q_2 \\ q_1 & -q_2 + \gamma_2(p_2 - 1) \end{pmatrix}^{-1} \begin{pmatrix} \lambda_1 \pi_1 \\ \lambda_2 \pi_2 \end{pmatrix}$$

and


$$\begin{pmatrix} m_{F,1} \\ m_{F,2} \end{pmatrix} = - \begin{pmatrix} -q_1 + \gamma_1((I-1)s_1 - 1) & q_2 \\ q_1 & -q_2 + \gamma_2((I-1)s_2 - 1) \end{pmatrix}^{-1} \cdot \begin{pmatrix} \gamma_1 r_1 & 0 \\ 0 & \gamma_2 r_2 \end{pmatrix} \begin{pmatrix} m_{L,1} \\ m_{L,2} \end{pmatrix}.$$

Using Proposition 3.4.3 we can determine under what condition these stationary means are well defined. A similar system of equations can be set up to determine the corresponding stationary (reduced) second moments, in self-evident notation denoted by


$$(v_{LL,1}, v_{LL,2}, v_{FF,1}, v_{FF,2}, v_{FF',1}, v_{FF',2}, v_{LF,1}, v_{LF,2});$$

as the underlying ideas are precisely the same as for the stationary means, we do not present the expressions here.

Importantly, our method is capable of providing near-instantaneous response. In our implementation we have used highly accurate, efficient and robust procedures that solve the systems of linear differential equations, available in any computational package. The alternative is Monte Carlo simulation, which is much slower: for instance generating the data points for Figure 3.1.b took about 20 minutes on a standard laptop, while Figure 3.2 (to be discussed in the next subsection) took roughly an hour. It is also noted that Monte Carlo provides just an *estimate* of the quantity under study, with the inherent uncertainty being quantified by means of a confidence interval. If one aims to have precise estimates, i.e., estimates with a narrow confidence interval, one has to bear in mind the rule of thumb that halving the width of the confidence interval requires that the number of sampled paths increases by a factor 4 (which follows by the central limit theorem).



(a) Transient means as numerical solution of differential equations



(b) Transient means approximated by simulations

Figure 3.1: Transient means of the leader's and follower's wealth. The numerical solutions have been evaluated by applying Euler's method. For the simulated approximation we have run 2000 simulations for a population of $I = 30$. The chosen parameters are: $q_{12} = 1/100$, $q_{21} = 5/100$, $\lambda_1 = 3$, $\lambda_2 = 1$, $\gamma_1 = 2$, $\gamma_2 = 1$, $p_1 = 0.3$, $p_2 = 0.6$, $1 - p_1 - (I-1)r_1 = 5/100$, $1 - p_2 - (I-1)r_2 = 1/100$, $1 - (I-1)s_1 = 5/100$ and $1 - (I-1)s_2 = 1/10$.

3.6.2 Poverty trap for single follower

Let f be the probability of the stationary wealth of an arbitrary follower, denoted by M_j for some follower $j = 2, \dots, I$, being below some critical threshold c ; this f can, in our stylized context, be considered as the probability of a follower ending up in the poverty trap. We propose a straightforward normal approximation, inspired by the central limit theorem (CLT), that is accurate if the scale of the underlying model is sufficiently high. A key property of such a CLT-based approximations is that only the first two moments of M_j play a role. In our case, this means

$$f = \mathbb{P}(M_j \leq c) \approx f_N := \Phi \left(\frac{c + 1/2 - m_F}{\sqrt{v_{FF}^\circ}} \right),$$

where the “ $+1/2$ ” is the standard continuity correction, and in addition

$$m_F := m_{F,1} + m_{F,2}, \quad v_{FF} := v_{FF,1} + v_{FF,2}, \quad v_{FF}^\circ := v_{FF} + m_F - m_F^2;$$

as usual, $\Phi(\cdot)$ denotes the cumulative distribution function of a standard normal random variable. To verify the accuracy of our findings, we have also estimated $\mathbb{P}(M_j \leq c)$ by performing 5000 independent simulation runs. The resulting distribution is indeed highly similar to the f_N resulting from the normal approximation, as shown in Figure 3.2.




Figure 3.2: Cumulative distribution function of the follower’s wealth, approximated through simulations and the normal approximation. For the simulated approximation we have run 5000 simulations for a population of $I = 30$. The chosen parameters are: $q_{12} = 1/100$, $q_{21} = 5/100$, $\lambda_1 = 10$, $\lambda_2 = 6$, $\gamma_1 = 4$, $\gamma_2 = 2$, $p_1 = 0.2$, $p_2 = 0.4$, $1 - p_1 - (I - 1)r_1 = 5/100$, $1 - p_2 - (I - 1)r_2 = 1/100$, $1 - (I - 1)s_1 = 3/100$ and $1 - (I - 1)s_2 = 7/100$.

Note that the followers do not operate independently, as they react to a common background process. As a consequence, the number of followers ending up in the state of a poverty trap, say B , is *not* binomially distributed (even though all of them experience the same probability f of doing so). In the remainder of this section we point out how to derive a good proxy for the distribution of B .

3.6.3 Poverty trap for full follower population

With B_j the indicator function of $M_j \leq c$, for $j = 2, \dots, I$, we are interested in approximating the distribution of

$$B := \sum_{j=2}^I B_j.$$

We propose a normal approximation, which relies on the central limit theorem, entailing that it will be particularly accurate as I grows. Define, for $j, j' = 2, \dots, I$ such that $j \neq j'$,

$$f' := \mathbb{P}(M_j \leq c, M_{j'} \leq c).$$

This probability can be approximated by its Gaussian counterpart. Applying continuity correction,

$$f' \approx f'_N := \mathbb{P}\left(M_j^\circ \leq c + \frac{1}{2}, M_{j'}^\circ \leq c + \frac{1}{2}\right),$$

with $(M_j^\circ, M_{j'}^\circ)$ bivariate normal with mean (m_F, m_F) and covariance matrix

$$\Sigma = \begin{pmatrix} v_{FF}^\circ & v_{FF'}^\circ \\ v_{FF'}^\circ & v_{FF}^\circ \end{pmatrix},$$

where $v_{FF}^\circ := v_{FF} + m_F - m_F^2$ (as before) and $v_{FF'}^\circ := v_{FF'} - m_F^2$. Notice that there are powerful numerical techniques to accurately evaluate bivariate normal probabilities; see for instance [40]. It now follows that

$$\begin{aligned} \mathbb{E} B &\approx \mu_B := (I-1)f_N, \\ \text{Var } B &\approx \sigma_B^2 := (I-1)f_N(1-f_N) + (I-1)(I-2)(f'_N - f_N^2). \end{aligned}$$

Applying a standard continuity correction, this gives rise to the following approximation: for $k = 1, \dots, I-1$,

$$\mathbb{P}(B = k) \approx \Phi\left(\frac{k + \frac{1}{2} - \mu_B}{\sigma_B}\right) - \Phi\left(\frac{k - \frac{1}{2} - \mu_B}{\sigma_B}\right).$$

As shown in Figure 3.3, the distribution following from the normal approximation matches the simulation-based approximation quite well.




Figure 3.3: Density function of the amount of follower's with wealth less or equal to 1, evaluated through simulations and the normal approximation. For the simulated approximation we have run 2000 simulations for a population of $I = 50$. The chosen parameters are: $q_{12} = 1/100$, $q_{21} = 5/100$, $\lambda_1 = 10$, $\lambda_2 = 6$, $\gamma_1 = 4$, $\gamma_2 = 2$, $p_1 = 0.2$, $p_2 = 0.4$, $1 - p_1 - (I-1)r_1 = 5/100$, $1 - p_2 - (I-1)r_2 = 1/100$, $1 - (I-1)s_1 = 3/100$ and $1 - (I-1)s_2 = 7/100$.

3.7 Application 2: opinion dynamics

In the field of opinion dynamics one studies, predominantly based on mathematical models, the evolution of opinions in a collection of agents. Arguably the most basic, yet meaningful model was proposed in [48], considering a set of n agents, each of them having an opinion on a particular subject. At every discrete time instant each agent updates her opinion based on the other agents' opinions. More specifically, the opinion of an agent's opinion at a certain point in time is a weighted sum of all agents' opinion at the previous time instant, with the weights summing to 1. To make the model more realistic, various extensions have been developed. In this context, one of the important contributions is [70], in which agents are assumed to have an internal opinion as well as an expressed opinion. In modelling terms, it means that the framework of [70] extends the one of [48] by dropping the assumption that the weights sum to 1 and by allowing (at any point in time) an external contribution to the opinion vector.

The goal of this section is to show that one can use the theory developed in the present chapter to make opinion dynamics models considerably more realistic. Compared to existing frameworks on opinion dynamics, two significant improvements can be made. In the first place, most of the existing models on opinion dynamics that allows for analytical investigation are of a deterministic nature, and as such do not incorporate the random effects present in a population of agents influencing each other's opinions; these can be included relying on our modelling framework. In the second place, it allows us, through the modulation mechanism, to incorporate randomly evolving external effects that have impact on the agents' opinions. Examples of external effects could relate to the current level of economic prosperity, or to the degree of access to digital communication, or to the geographic area agents live in; these variables change over time (where the natural timescale could correspond to years) which result in different opinion dynamics.

Let us consider the first advantage, i.e., the option of introducing stochasticity, in more detail. The framework that we developed can in fact be interpreted as a stochastic generalisation of the model studied in [70], which (in its most elementary form) can be summarised as follows. Let Y_t a vector of opinions, W_t a matrix that describes the effects of each opinion held at time $t - 1$ on the opinions held at time t , X_t a matrix of scores on exogenous variables, B_t a vector of coefficients giving the effects of each of the exogenous variables, α_t a scalar weight corresponding to the endogenous conditions and β_t a scalar weight corresponding to the exogenous conditions. Then the recurrence relation that defines the model in [70] is given, for $t \in \mathbb{N}$, by

$$Y_t = \alpha_t W_t Y_{t-1} + \beta_t X_t B_t.$$

The relation with our model can be seen immediately by comparing this recurrence relation with the differential equation of the transient means in Proposition 3.4.1. Indeed, just as in [70], the opinion of an agent is construed as a linear combination of its own and the others' opinions before the time of redistribution. An important additional advantage of our approach is the option to explicitly characterise all (reduced) moments of the opinion vector through systems of differential equations, by repeated differentiation of the relation featuring in Proposition 3.3.1. To our best knowledge, this is a novel analytical tool that was not provided in other stochastic generalisations of the Friedkin-Johnsen model [70].

3.7.1 Opinion dynamics with Markovian background process

In this example we consider opinion dynamics in a two-group population (groups A and B) where individuals of one group – say group A – have an opinion update mechanism that alternates between a “normal mode” and an “adapted mode”. This framework may be interpreted as a stylised model for groups of individuals occasionally visiting larger events as conferences, demonstrations, or political gatherings. Formally we describe the model as follows.

- The background process $X(t)$ has two states, i.e., one corresponding to the normal mode (corresponding to state 1) and one to the adapted mode (corresponding to state 2).

- We consider the situation that individuals do not automatically increase the value of their opinions when the background process is in the normal mode. Since group B agents are not affected by the background process, we thus have $\lambda_{A,1} = \lambda_{B,1} = \lambda_{B,2} = 0$. Additionally, we let group A agents automatically strengthen their opinion with rate $\lambda_{A,2} \geq 0$ in the adapted mode, i.e., they only strengthen their opinion as a result of external influence by attending conference, political meetings, etc.

- We let the Poisson rate of opinion redistribution $\gamma > 0$ be unaffected by the background process. Since group A separates itself from group B when the background process is in the adapted state, we assume that group A and B individuals do not affect each other’s opinion in that state. In addition, in this example we allow group A agents to be “overenthusiastic” in the adapted state, resulting in distributing more opinion units than they originally have.

We first consider the situation that the background process is in state 1. Let there be I_A agents in group A, and I_B in group B. For $j \in \{1, \dots, I_A\}$ (i.e., corresponding to agents in group A), we let the vector

$$(W_{j,1,1}, \dots, W_{j,I_A,1}, W_{j,I_A+1,1}, \dots, W_{j,I_A+I_B,1})$$

be multinomially distributed with parameters 1 and $(p_{AA,1}, \dots, p_{AA,1}, p_{AB,1}, \dots, p_{AB,1})$; for $j \in \{I_A + 1, \dots, I_A + I_B\}$ (i.e., corresponding to agents in group B) these parameters are 1 and $(p_{BA,1}, \dots, p_{BA,1}, p_{BB,1}, \dots, p_{BB,1})$. Here we consider the situation that

$$I_A p_{AA,1} + I_B p_{AB,1} = 1 \quad \text{and} \quad I_A p_{BA,1} + I_B p_{BB,1} = 1,$$

effectively meaning that “opinion mass cannot leak away from the system” when the background process is in state 1.

When the background process is in state 2, we let for $j \in \{I_A + 1, \dots, I_A + I_B\}$ (i.e., for agents of group B)

$$(W_{j,1,2}, \dots, W_{j,I_A,2}, W_{j,I_A+1,2}, \dots, W_{j,I_A+I_B,2})$$

be multinomially distributed with parameters 1 and $(0, \dots, 0, p_{BB,2}, \dots, p_{BB,2})$. Here, group B opinion does not leak away: $I_B p_{BB,2} = 1$. For $j \in \{1, \dots, I_A\}$ (i.e., for agents of group A) we introduce an additional mechanism: with probability $0 \leq \alpha \leq 1$ the unit of opinion to be distributed becomes two units and with probability $1 - \alpha$ it remains one unit. Afterwards, each unit is multinomially distributed with parameters 1 and $(p_{AA,2}, \dots, p_{AA,2}, 0, \dots, 0)$, with $I_A p_{AA,2} = 1$.

We would like to stress that the mechanism described above is just an example of one specific type of opinion dynamics we can handle. Various alternative models can be dealt similarly (more than two groups, more than two background states, the “ α -jumps” corresponding with multiplying the “opinion unit” with a number different from two, etc.).

The model proposed allows groups of individuals of group A to “create opinion” when the background state is 2 (by $\lambda_{A,2} > 0$), and to multiply each of their existing opinion units by 2 with probability α . Since “opinion mass” does not leave the population in this setup, we typically expect opinions to grow unboundedly in time when $\lambda_{A,2} > 0$ or $\alpha > 0$. In the next section we show that the aforementioned is indeed true for the means of the agent’s opinion.

3.7.2 Means and variances

For ease, we assume that the background process is in stationarity at time 0. With the results found in Section 3.4.3 and the notation as in Proposition 3.4.1, we have the following collection of differential equations describing the transient means:

$$\mathbf{m}'(t) = A \mathbf{m}(t) + \Lambda \boldsymbol{\pi},$$

with $\mathbf{m}(t) = (m_{A,1}(t), m_{A,2}(t), m_{B,1}(t), m_{B,2}(t))^T$,

$$A = \begin{pmatrix} -q_{12} + \gamma(I_A p_{AA,1} - 1) & q_{21} & \gamma I_B p_{BA,1} & 0 \\ q_{12} & -q_{21} + \gamma \alpha & 0 & 0 \\ \gamma I_A p_{AB,1} & 0 & -q_{12} + \gamma(I_B p_{BB,1} - 1) & q_{21} \\ 0 & 0 & q_{12} & -q_{21} \end{pmatrix}$$

and $\Lambda = \text{diag}\{0, \lambda_{A,2}, 0, 0\}$. For the transient second moments, similar differential equations can be derived as well, relying on the results presented in Section 3.5.1.

It is straightforward to verify that a necessary condition for the existence of a steady state is $\alpha = 0$ and $\lambda_{A,2} = 0$, as anticipated in the previous section. Under these conditions, opinions are only redistributed since no additional opinion is “created” and, by assumption, no opinion is “lost”. The steady state of this model is, by solving $A \mathbf{m} = \mathbf{0}$, proportional to the vector

$$\left(1, \frac{q_{12}}{q_{21}}, \frac{p_{AB,1}}{p_{BA,1}}, \frac{p_{AB,1} q_{12}}{p_{BA,1} q_{21}}\right)^T. \quad (3.4)$$


This vector should be normalised such that the total “opinion mass” equals its initial value (i.e., $\mathbf{1}^T \mathbf{M}(0)$). Indeed, if the setup of the system corresponds to spending longer periods of time in state 2, or equivalently $q_{12} > q_{21}$, then we expect $m_{A,2} > m_{A,1}$. Also, when group A agents distribute their opinions with a higher probability to group B agents than the probability of group B agents distributing their opinion to group A agents, thus $p_{AB,1} > p_{BA,1}$, we expect $m_{B,1} > m_{A,1}$. Both properties are in line with the steady-state expression (3.4).

When $\alpha = 0$ and $\lambda_{A,2} = 0$, the adapted state can be interpreted as a temporary interruption of the opinion formation process between groups A and B; in the adapted state, agents of both groups do not distribute their opinion to other group’s agents. This interruption should not affect the steady state of the system. Indeed, the steady states of the groups

are related through $m_A = (p_{BA,1}/p_{AB,1}) m_B$ which does not involve the background process parameters q_{12} and q_{21} .

Figure 3.4 shows that the simulation-based approximation of the means aligns with the numerical solution of the differential equation of the means. In this setup the abscissa ω of the matrix A is equal to 0, thus not satisfying the stability condition of Proposition 3.4.3. We see that, although $\omega = 0$, the transient means do converge to a steady state. Note that this does not contradict Proposition 3.4.3: the condition $\omega < 0$ is a *sufficient* condition, hence there can be stability even though $\omega < 0$ is not fulfilled. In Section 3.7.3 we provide an example where there is no stability while the abscissa ω is 0. In the setup of Figure 3.4, the agents of group A tend to give more attention to the topic of interest than the agents of group B, in the sense that $I_A p_{AA,1} + I_B p_{BA,1} = 5/3$ while $I_A p_{AB,1} + I_B p_{BB,1} = 8/9$. This “unequal attentiveness” typically results in polarisation, as discussed in [35].

In Figure 3.5 we show the effect of different background parameters q_{12} and q_{21} . The curve of the transient means corresponding to spending a longer fraction of time in the adapted state (i.e., the one with the higher q_{12}) moves more slowly to the steady state.



(a) Transient means as numerical solution of differential equations

(b) Transient means approximated by simulations

Figure 3.4: Transient means of group A's and B's opinion. The numerical solutions have been evaluated by applying Euler's method. For the simulated approximation we have run 2000 simulations. The group sizes are $I_A = 10$ and $I_B = 30$. The initial opinions of group A and B agents are 1 and 5 respectively. The aggregate opinion of all agents, $I_A m_A(t) + I_B m_B(t)$, remains constant and equals $I_A m_A(0) + I_B m_B(0) = 160$, as we have enforced “conservation of opinion”. The chosen parameters are: $q_{12} = 3/10$, $q_{21} = 2/10$, $\lambda_{A,2} = 0$, $\alpha = 0$, $\gamma = 5/8$, $I_A p_{AA,1}/I_B p_{AB,1} = 2$ and $I_A p_{BA,1}/I_B p_{BB,1} = 4/5$. The abscissa ω is 0.

3.7.3 Unbounded opinion mass and relative opinions

We now consider the setup where α or $\lambda_{A,2}$ is assumed to be strictly positive. As reasoned earlier, in this situation we expect the total “opinion mass” to grow beyond any bound. In Figure 3.6a we show the evolution of the transient means in the situations $(\alpha, \lambda_{A,2}) = (0, 2)$ and $(\alpha, \lambda_{A,2}) = (1/10, 0)$.

It is noted that, in principle, unbounded opinions (in our case due to $\alpha > 0$ or $\lambda_{A,2} > 0$) are an outcome of the model that cannot be supported by empirical data.

However, various techniques have been proposed in order to avert the phenomenon of unbounded opinions while maintaining the mentioned interaction mechanism, as explained in [66]. One of the ideas is to interpret the modelled opinions as *relative* opinions [35].





Figure 3.5: Effect of a longer period of time in the adapted state on the transient means. The transient means have been evaluated by applying Euler’s method. The longer period of time in the adapted state is obtained by increasing q_{12} (i.e., we now set $q_{12} = 1$, versus $q_{12} = 3/10$ in the base case). All the other parameters are identical to the parameters used in Figure 3.4. The abscissa ω is 0 in both situations.

This means that the absolute values of the entries of the opinion vector have no meaning, but that interpretation is only given to the fractions between these entries. Concretely, the i -th entry divided by the j -th entry reflects the relative opinion of agent i with respect to agent j .

In Figure 3.6b it is shown how the unbounded growth of opinion mass in Figure 3.6a changes into a stabilising growth when considering the relative perspective. This illustrates that, when dealing with instances of the stochastic model where opinion mass may grow unboundedly, one may consider working with relative opinions instead. With this interpretation, one can still obtain insight into the effects that the specific interaction mechanism between the agents has on the opinion formation process of the entire population. Further empirical and technical arguments in favour of working with relative opinions, rather than their absolute counterparts, are discussed in great detail in [35].



(a) Transient means “growing” without limits

(b) Transient means interpreted as relative opinions

Figure 3.6: Unbounded transient means become bounded when interpreted as relative opinions. The numerical solutions have been evaluated by applying Euler’s method. All the parameters chosen are identical to the setup as in Figure 3.4 except for $\lambda_{A,2}$ and α . The abscissa ω is 0 in the setup with $\lambda_{A,2} = 2$. The abscissa ω is 0.027 in the setup with $\alpha = 10\%$.

3.8 Application 3: file storage systems

Consider a system in which users generate files. For safety reasons, the files that have been saved at the clients' locations are periodically copied to a central storage location where a backup is made. The frequency of copying the clients' files has to be sufficiently high to make sure that a relatively small amount has no centrally stored duplicate. In this section we discuss a sequence of models describing the dynamics of such file storage systems, starting with the most rudimentary variant. For more background on various aspects of data storage networks, we refer to e.g. [130].

3.8.1 Basic variant

A stylised model that describes the dynamics of this system is the following. Let $M_1(t)$ record the number of files at the clients' devices that have not been copied to the central storage unit by time $t \geq 0$, and let $M_2(t)$ the number of files at the storage unit at time $t \geq 0$. Let λ be the Poisson rate at which the aggregate client population generates files. Let the times between subsequent backups be exponentially distributed with mean γ^{-1} . In this system without modulation, we observe that $W_{12} \equiv W_{22} \equiv 1$, and $W_{11} \equiv W_{21} \equiv 0$. It can be verified from the results of Section 3.4 that

$$m'_1(t) = \lambda - \gamma m_1(t), \quad m'_2(t) = \gamma m_1(t).$$

It is not hard to verify that the first moments can be computed explicitly: if the system starts empty, for $t \geq 0$,

$$m_1(t) = \frac{\lambda}{\gamma}(1 - e^{-\gamma t}), \quad m_2(t) = \lambda t - \frac{\lambda}{\gamma}(1 - e^{-\gamma t}).$$

These expressions can be used to determine the optimal backup rate γ^* , for instance by considering a cost function that encompasses the per-backup cost (with proportionality constant $\kappa_B > 0$) and the cost associated with the number of files that have not been copied yet (with proportionality constant $\kappa_{NC} > 0$). Concretely, for given time horizon $t \geq 0$, this leads to the optimisation problem

$$\min_{\gamma > 0} F_1(\gamma) + F_2(\gamma), \quad \text{with } F_1(\gamma) := \gamma t \kappa_B, \quad F_2(\gamma) := \frac{\lambda}{\gamma}(1 - e^{-\gamma t}) \kappa_{NC},$$

with the objective function $F(\gamma) := F_1(\gamma) + F_2(\gamma)$ being convex. It takes some calculus to verify that if $F'(0) \geq 0$, or (equivalently) $2\kappa_B \geq \lambda t \kappa_{NC}$, then $\gamma = \gamma^* = 0$ is optimal, which in practical terms means that one should very frequently update; if on the other hand $F'(0) < 0$, or (equivalently) $2\kappa_B < \lambda t \kappa_{NC}$, there is a strictly positive optimal update rate γ^* . In the latter case, γ^* cannot be computed in closed form, but that can be easily numerically evaluated; see also Figure 3.7

We can also find the (reduced) second moments $v_{ij}(t)$, with $i, j = 1, 2$, by using the techniques from Section 3.5. Indeed, for $t \geq 0$, we have

$$\begin{aligned} v'_{11}(t) &= 2\lambda m_1(t) - \gamma v_{11}(t), \\ v'_{12}(t) &= \lambda m_2(t) - \gamma v_{12}(t), \\ v'_{22}(t) &= \gamma v_{11}(t) + 2\gamma v_{12}(t). \end{aligned}$$




Figure 3.7: The functions $F(\gamma)$, $F_1(\gamma)$, and $F_2(\gamma)$ for $2\kappa_B \geq \lambda t \kappa_{NC}$ (left panel) and for $2\kappa_B < \lambda t \kappa_{NC}$ (right panel). In the former case $\gamma^* = 0$, whereas in the latter case $\gamma^* > 0$.

Solving the differential equations we obtain

$$\begin{aligned} v_{11}(t) &= \left(\frac{\lambda}{\gamma}\right)^2 \left[2(1 - e^{-\gamma t}) - 2\gamma t e^{-\gamma t} \right] \\ v_{12}(t) &= \left(\frac{\lambda}{\gamma}\right)^2 \left[\gamma t (1 + e^{-\gamma t}) + 2(e^{-\gamma t} - 1) \right] \\ v_{22}(t) &= \left(\frac{\lambda}{\gamma}\right)^2 \left[2(1 - e^{-\gamma t}) + \gamma^2 t^2 - 2\gamma t \right]. \end{aligned}$$

Using the first and (reduced) second moments, we can derive the following expressions for the variances:

$$\begin{aligned} \mathbb{V}\text{ar } M_1(t) &= \left(\frac{\lambda}{\gamma}\right)^2 \left[-2\gamma t e^{-\gamma t} - e^{-2\gamma t} - \frac{\gamma}{\lambda} e^{-\gamma t} + 1 + \frac{\gamma}{\lambda} \right] \\ \mathbb{V}\text{ar } M_2(t) &= \left(\frac{\lambda}{\gamma}\right)^2 \left[-2\gamma t e^{-\gamma t} - e^{-2\gamma t} + \frac{\gamma}{\lambda} e^{-\gamma t} + \frac{\gamma^2}{\lambda} t + 1 - \frac{\gamma}{\lambda} \right]. \end{aligned}$$

3.8.2 Faulty upload link

This model can be further refined in various ways. In the first place, the link between the clients and the central storage may be faulty, in that it alternates between “functioning” and “broken”. Suppose that the time the link is up (down, respectively) is exponentially distributed with parameter q_U (q_D , respectively). Then it easily seen, using the techniques developed in Section 3.5, that, in self-evident notation, with the second index corresponding to the link being up or down,

$$\begin{aligned} m'_{1U}(t) &= -q_U m_{1U}(t) + q_D m_{1D}(t) + \lambda \pi_U(t) - \gamma m_{1U}(t), \\ m'_{1D}(t) &= q_U m_{1U}(t) - q_D m_{1D}(t) + \lambda \pi_D(t), \\ m'_{2U}(t) &= -q_U m_{2U}(t) + q_D m_{2D}(t) + \gamma m_{1U}(t), \\ m'_{2D}(t) &= q_U m_{2U}(t) - q_D m_{2D}(t); \end{aligned}$$

here the probabilities $\pi_U(t)$ and $\pi_D(t)$ follow from

$$\pi_U(t) = \pi_U(0)\pi_{UU}(t) + \pi_D(0)\pi_{DU}(t), \quad \pi_D(t) = \pi_U(0)\pi_{UD}(t) + \pi_D(0)\pi_{DD}(t),$$

with, abbreviating $q := q_D + q_U$,

$$\pi_{DU}(t) = 1 - \pi_{DD}(t) = \frac{q_D}{q}(1 - e^{-qt}), \quad \pi_{UD}(t) = 1 - \pi_{UU}(t) = \frac{q_U}{q}(1 - e^{-qt}).$$

The (reduced) second moments can be found relying on the theory of Section 3.5. As these can be found using the same line of reasoning as for the first moments, we do not provide the resulting differential equations here.

3.8.3 Failures in storage unit

A second extension includes the effect that the storage system itself can also be faulty, where upon failure all stored files are lost. We let these failure occur according to a Poisson process with intensity $\bar{\gamma}$. Let $M_3(t)$ denote the number of lost files. While the differential equations for $m_{1U}(t)$ and $m_{1D}(t)$ remain unchanged, the other ones become

$$\begin{aligned} m'_{2U}(t) &= -q_U m_{2U}(t) + q_D m_{2D}(t) + \gamma m_{1U}(t) - \bar{\gamma} m_{2U}(t), \\ m'_{2D}(t) &= q_U m_{2U}(t) - q_D m_{2D}(t) - \bar{\gamma} m_{2D}(t); \\ m'_{3U}(t) &= -q_U m_{3U}(t) + q_D m_{3D}(t) + \bar{\gamma} m_{2U}(t), \\ m'_{3D}(t) &= q_U m_{3U}(t) - q_D m_{3D}(t) + \bar{\gamma} m_{2D}(t). \end{aligned}$$

For a company offering this service, a typical design goal would be: how small should q_U , respectively $\bar{\gamma}$, be to make sure that the mean number of lost files at time t , i.e., $m_{3U}(t) + m_{3D}(t)$, is below a given threshold? More advanced criteria could also involve the corresponding variance.

To mitigate the effect of file loss due to failures of the storage unit, an evident policy is to copy all files to *multiple* storage units. For instance, one could consider the mechanism in which every file is simultaneously copied to two storage units at Poisson instants (again with rate γ), and that each of the storage units fails at Poisson instants (again with rate $\bar{\gamma}$). This means that a file is lost only if both storage units have failed. The above analysis (for the case of a single storage unit, that is) can be extended in an evident manner to this situation, thus facilitating the quantification of the performance gain due to making multiple backups. In order to decide whether such a policy should be implemented, this gain should be compared to the cost of the additional storage unit.

3.9 Concluding remarks

This chapter developed a versatile Markovian model that describes the dissemination of wealth over a population, but with the potential to be broadly applied across a wide range of disciplines. It is demonstrated that the evaluation of transient moments reduces to solving a system of coupled linear differential equations, while their stationary counterparts require solving (ordinary) linear systems of equations. Also a stability condition has been developed. Three examples evidence the model's wide application potential.

In various directions there is scope for generalisations. It is in particular noted that, conditional on the state of the background process, the wealth units evolve independently, which is not in all application areas a realistic mechanism. In addition, one may try to relax the various exponentiality assumptions.